

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	25192967	@ad<"20031126"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 12:35
L2	2	1 and "systolic memory arrays"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 12:37
L3	0	1 and (pipelin\$3 pipe-lin\$3) with array same read same write same bank same systolic	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 14:02
L4	29	1 and (pipelin\$3 pipe-lin\$3) with array same read same write same bank	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 12:46
L5	7	4 and (pipelin\$3 pipe-lin\$3) with registers	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 12:47
L6	1	1 and (pipelin\$3 pipe-lin\$3) with array same read same write same bank and systolic	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 14:07
L8	1	1 and (pipelin\$3 pipe-lin\$3) with banks same read\$3 with writ\$3 and systolic	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 14:12

EAST Search History

L9	1	1 and (pipelin\$3 pipe-lin\$3) same arrays same banks same read\$3 with writ\$3 and systolic	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 14:14
L10	58	1 and (pipelin\$3 pipe-lin\$3) same arrays same banks same read\$3 with writ\$3	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 14:20
L11	0	1 and (pipelin\$3 pipe-lin\$3) same arrays same banks same read\$3 with writ\$3 and "address pipe"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 14:22
L12	39	1 and (pipelin\$3 pipe-lin\$3) same arrays same banks same read\$3 with writ\$3 and (begin\$4 start\$3 first\$2 initializ\$5 commenc\$5) with bank	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 14:48
L13	5552	(lu.in. somasekhar.in. ye.in.)	US-PGPUB	OR	ON	2007/03/26 14:50
L14	8	13 and systolic.clm.	US-PGPUB	OR	ON	2007/03/26 14:50
L15	1	13 and systolic.clm. and banks.clm.	US-PGPUB	OR	ON	2007/03/26 14:50
L16	38418	(lu.in. somasekhar.in. ye.in.)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 14:50
L17	1	16 and systolic.clm. and banks.clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 14:51

EAST Search History

L18	661	711/169.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 15:09
L19	1771	711/167.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 15:10
L20	1561	711/100.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 15:10
L21	271	711/101.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 15:10
L22	2564	711/170.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 15:13
L23	2	"20050114618".pn.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 19:39
L24	957	710/36.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 17:14

EAST Search History

L25	580	710/38.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 18:35
L26	8	(18 19 20 21 22 24 25) and (2 4 5 6 8 9 10 12)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 18:43
L28	479	710/3.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 18:41
L29	3534	365/189.01.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 18:42
L30	4078	365/230.03.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 18:42
L31	18	(18 19 20 21 22 24 25 28 29 30) and (2 4 5 6 8 9 10 12)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 19:07
L32	20	1 and (pipelin\$3 pipe-lin\$3) near3 registers same arrays same banks	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 19:09

EAST Search History

L33	1	23 and "one side"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/26 19:39
-----	---	-------------------	--	----	----	------------------



Welcome United States Patent and Trademark Office

[Search Session History](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Mon, 26 Mar 2007, 7:53:30 PM EST

Edit an existing query or
compose a new query in the
Search Query Display.

Search Query Display

Select a search number (#)
to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Recent Search Queries

- #1 ((systolic<in>metadata) <and> (memory arrays<in>metadata))<and> (pipelined<in>metadata)
- #2 ((systolic<in>metadata) <and> (pipelined<in>metadata))
- #3 ((systolic<in>metadata) <and> (pipelined<in>metadata))
- #4 ((systolic<in>metadata) <and> (pipelined<in>metadata))
- #5 ((systolic<in>metadata) <and> (pipelined registers<in>metadata))
- #6 (pipelined registers<in>metadata)
- #7 ((systolic memory arrays<in>metadata) <and> (pipeline registers<in>metadata))
- #8 ((systolic<in>metadata) <and> (pipeline registers<in>metadata))<and> (banks<in>metadata)
- #9 ((systolic<in>metadata) <and> (banks<in>metadata))<and> (pipeline<in>metadata)
- #10 ((pipelined<in>metadata) <and> (systolic<in>metadata))<and> (arrays<in>metadata)
- #11 ((pipelined<in>metadata) <and> (systolic<in>metadata))<and> (arrays<in>metadata)
- #12 ((systolic memory arrays<in>metadata) <and> (pipeline<in>metadata))

Indexed by
 Inspect°

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

[systolic](#) [arrays](#) [banks](#) [pipeline](#) [read](#) [write](#) [address](#) [registers](#) [pipes](#)

Found 42 of 198,991

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 42

Result page: [1](#) [2](#) [3](#) [next](#)Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Warp architecture and implementation](#)



M. Annaratone, E. Arnould, T. Gross, H. T. Kung, M. S. Lam

 June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th annual international symposium on Computer architecture ISCA '86**, Volume 14 Issue 2

Publisher: IEEE Computer Society Press, ACM Press

Full text available: pdf(1.17 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes the scan line array processor (SLAP), a new architecture designed for high-performance yet low-cost image computation. A SLAP is a SIMD linear array of processors, and hence is easy to build and scales well with VLSI technology; yet appropriate special features and programming techniques make it efficient for a surprisingly wide variety of low and medium level computer vision tasks. We describe the basic SLAP concept and some of its variants, discuss a particular planne ...

2 [The family of concurrent logic programming languages](#)



Ehud Shapiro

 September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

Publisher: ACM Press

Full text available: pdf(9.62 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Concurrent logic languages are high-level programming languages for parallel and distributed systems that offer a wide range of both known and novel concurrent programming techniques. Being logic programming languages, they preserve many advantages of the abstract logic programming model, including the logical reading of programs and computations, the convenience of representing data structures with logical terms and manipulating them using unification, and the amenability to metaprogrammin ...

3 [Warp architecture and implementation](#)




Marco Annaratone, Emmanuel Arnould, Thomas Gross, H. T. Kung, Monica S. Lam, Onat Menzilcioglu, Ken Sarocky, Jon A. Webb

 August 1998 **25 years of the international symposia on Computer architecture (selected papers) ISCA '98**

Publisher: ACM Press

Full text available:

Additional Information:



 [pdf\(1.17 MB\)](#)[full citation](#), [references](#), [index terms](#)

4 [A decade of reconfigurable computing: a visionary retrospective](#)

R. Hartenstein

March 2001 **Proceedings of the conference on Design, automation and test in Europe DATE '01****Publisher:** IEEE PressFull text available:  [pdf\(768.00 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

5 [Overview of a high-performance programmable pipeline structure](#)


 Franc_ois Bodin, Franc_ois Charot, Charles WagnerJune 1989 **Proceedings of the 3rd international conference on Supercomputing ICS '89****Publisher:** ACM PressFull text available:  [pdf\(2.05 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper aims at describing a high-performance programmable pipeline architecture consisting of a linear array of PCS processors. The PCS processor which is capable of performing 20 million floating-point operations per second (20 MFLOPS) has been built from off-the-shelf chips on a wire-wrapped board. The prototype processor is attached to a SUN-3 workstation. Efficient microcode is generated using the microcode compiler that has been designed and implemented. The microcode op ...

6 [Three-dimensional finite-element analyses: implications for computer architectures](#)

 Valerie E. Taylor, Abhiram Ranade, David G. MesserschmittAugust 1991 **Proceedings of the 1991 ACM/IEEE conference on Supercomputing Supercomputing '91****Publisher:** ACM PressFull text available:  [pdf\(849.04 KB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

7 [Evaluation of the Raw Microprocessor: An Exposed-Wire-Delay Architecture for ILP and Streams](#)

 Michael Bedford Taylor, Walter Lee, Jason Miller, David Wentzlaff, Ian Bratt, Ben Greenwald, Henry Hoffmann, Paul Johnson, Jason Kim, James Psota, Arvind Saraf, Nathan Shnidman, Volker Strumpfen, Matt Frank, Saman Amarasinghe, Anant AgarwalMarch 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04**, Volume 32 Issue 2**Publisher:** IEEE Computer Society, ACM PressFull text available:  [pdf\(376.05 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#)

This paper evaluates the Raw microprocessor. Raw addresses the challenge of building a general-purpose architecture that performs well on a larger class of stream and embedded computing applications than existing microprocessors, while still running existing ILP-based sequential programs with reasonable performance in the face of increasing wire delays. Raw approaches this challenge by implementing plenty of on-chip resources - including logic, wires, and pins - in a tiled arrangement, and exposing the ...

8 [The white dwarf: a high-performance application-specific processor](#)

A. Wolfe, M. Breternitz, C. Stephens, A. L. Ting, D. B. Kirk, R. P. Bianchini, J. P. Shen



May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual International Symposium on Computer architecture ISCA '88**,
Volume 16 Issue 2

Publisher: IEEE Computer Society Press, ACM Press

Full text available: pdf(1.40 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents the design and implementation of a high-performance special-purpose processor, called The White Dwarf, for accelerating finite element analysis algorithms. The White Dwarf CPU contains two Am29325 32-bit floating-point processors and one Am29332 32-bit ALU, and employs a wide-instruction word architecture in which the application algorithm is directly implemented in microcode. The entire system is VME-bus compatible and interfaces with a SUN 31160 host. The syste ...

9 [A hardware accelerator for maze routing](#)



Y. Won, S. Sahni, Y. El-ziq

October 1987 **Proceedings of the 24th ACM/IEEE conference on Design automation DAC '87**

Publisher: ACM Press

Full text available: pdf(871.73 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A hardware accelerator for the maze routing problem is developed. This accelerator consists of three 3 stage pipelines. Banked memory is used to avoid memory read/write conflicts and obtain maximum efficiency.

10 [Computing multi-colored polygonal masks in pipeline architecture and its application to automated visual inspection](#)



Jorge L. C. Sanz, Its'hak Dinstein, Dragutin Petkovic

April 1987 **Communications of the ACM**, Volume 30 Issue 4

Publisher: ACM Press

Full text available: pdf(2.56 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

New techniques for computing multicolored polygonal masks for image analysis and computer vision applications are presented. The procedures do not require random access of the image memory. They are based on efficient generation of coordinate-reference images (ramps) and other simple general purpose architectural features such as look-up tables. The techniques presented are, unlike their predecessors, highly parallel and can be efficiently implemented in existing pipeline image processors. ...

11 [A high-speed network interface for distributed-memory systems: architecture and applications](#)



Peter Steenkiste

February 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 1

Publisher: ACM Press

Full text available: pdf(993.12 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

Distributed-memory systems have traditionally had great difficulty performing network I/O at rates proportional to their computational power. The problem is that the network interface has to support network I/O for a supercomputer, using computational and memory bandwidth resources similar to those of a workstation. As a result, the network interface becomes a bottleneck. In this article we present an I/O architecture that addresses these problems and supports high-speed network I/O on dist ...

Keywords: I/O architecture, application-managed I/O, data reshuffling, distributed memory systems, network interface, outboard buffering, protocol processing, resource management

12 Using Lookahead to reduce memory bank contention for decoupled operand references



Peter L. Bird, Richard A. Uhlig

August 1991 **Proceedings of the 1991 ACM/IEEE conference on Supercomputing Supercomputing '91**

Publisher: ACM Press

Full text available: pdf(1.09 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

13 Accelerators: Automatic mapping of nested loops to FPGAS



Uday Bondhugula, J. Ramanujam, P. Sadayappan

March 2007 **Proceedings of the 12th ACM SIGPLAN symposium on Principles and practice of parallel programming PPOPP '07**

Publisher: ACM Press

Full text available: pdf(245.06 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper presents a framework for automatic mapping of perfectly nested loops with constant dependences onto regular processor arrays, suitable for direct implementation on Field Programmable Gate Arrays (FPGAs). The problem is modeled as that of finding a suitable completion procedure for a full-rank linear transformation on the iteration space. The approach enables extraction of necessary degrees of communication-free and pipelined parallelism to optimize performance under the resource constraints.

Keywords: FPGA, FPGA compilation, control signals, linear transformation, nested loops, regular processor arrays, resource constraints, scheduling

14 GPMB—software pipelining branch-intensive loops

Zhihong Tang, Gang Chen, Chihong Zhang, Yingwei Zhang, Bogong Su, Stanley Habib

December 1993 **Proceedings of the 26th annual international symposium on Microarchitecture MICRO 26**

Publisher: IEEE Computer Society Press

Full text available: pdf(906.47 KB) Additional Information: [full citation](#), [references](#), [citations](#)

Keywords: branch overlapping, branch-intensive loop-level parallelism, multi-branch switch, processing element

15 Trident: a scalable architecture for scalar, vector, and matrix operations

Mostafa I. Soliman, Stanislav G. Sedukhin

January 2002 **Australian Computer Science Communications, Proceedings of the seventh Asia-Pacific conference on Computer systems architecture CRPIT '02**, Volume 24 Issue 3

Publisher: Australian Computer Society, Inc., IEEE Computer Society Press

Full text available: pdf(814.51 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Within a few years it will be possible to integrate a billion transistors on a single chip. At this integration level, we propose using a high level ISA to express parallelism to hardware instead of using a huge transistor budget to dynamically extract it. Since the fundamental data structures for a wide variety of applications are scalar, vector, and matrix, our proposed Trident processor extends the classical vector ISA with matrix

operations. The Trident processor consists of a set of parallel ...

Keywords: data parallelism, parallel processing, ring register file, scalable hardware, vector/matrix processing

16 Join processing in relational databases



Priti Mishra, Margaret H. Eich

March 1992 **ACM Computing Surveys (CSUR)**, Volume 24 Issue 1

Publisher: ACM Press

Full text available: [pdf\(4.42 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The join operation is one of the fundamental relational database query operations. It facilitates the retrieval of information from two different relations based on a Cartesian product of the two relations. The join is one of the most difficult operations to implement efficiently, as no predefined links between relations are required to exist (as they are with network and hierarchical systems). The join is the only relational algebra operation that allows the combining of related tuples from ...

Keywords: database machines, distributed processing, join, parallel processing, relational algebra

17 Reconfigurable computing: a survey of systems and software



Katherine Compton, Scott Hauck

June 2002 **ACM Computing Surveys (CSUR)**, Volume 34 Issue 2

Publisher: ACM Press

Full text available: [pdf\(710.56 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Due to its potential to greatly accelerate a wide variety of applications, reconfigurable computing has become a subject of a great deal of research. Its key feature is the ability to perform computations in hardware to increase performance, while retaining much of the flexibility of a software solution. In this survey, we explore the hardware aspects of reconfigurable computing machines, from single chip architectures to multi-chip systems, including internal structures and external coupling. We ...

Keywords: Automatic design, FPGA, field-programmable, manual design, reconfigurable architectures, reconfigurable computing, reconfigurable systems

18 A review of vessel extraction techniques and algorithms



Cemil Kirbas, Francis Quek

June 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 2

Publisher: ACM Press

Full text available: [pdf\(8.06 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Vessel segmentation algorithms are the critical components of circulatory blood vessel analysis systems. We present a survey of vessel extraction techniques and algorithms. We put the various vessel extraction approaches and techniques in perspective by means of a classification of the existing research. While we have mainly targeted the extraction of blood vessels, neurovascular structure in particular, we have also reviewed some of the segmentation methods for the tubular objects that show ...

Keywords: Magnetic resonance angiography, X-ray angiography, medical imaging,

neurovascular, vessel extraction

19 Splash 2



Jeffrey M. Arnold, Duncan A. Buell, Elaine G. Davis

June 1992 **Proceedings of the fourth annual ACM symposium on Parallel algorithms and architectures SPAA '92**

Publisher: ACM Press

Full text available: pdf(719.35 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

20 Input data reuse in compiling window operations onto reconfigurable hardware



Zhi Guo, Betul Buyukkurt, Walid Najjar

June 2004 **ACM SIGPLAN Notices , Proceedings of the 2004 ACM SIGPLAN/SIGBED conference on Languages, compilers, and tools for embedded systems LCTES '04**, Volume 39 Issue 7

Publisher: ACM Press

Full text available: pdf(253.01 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Balancing computation with I/O has been considered as a critical factor of the overall performance for embedded systems in general and reconfigurable computing systems in particular. Data I/O often dominates the overall computation performance for window operation, which are frequently used in image processing, image compression, pattern recognition and digital signal processing. This problem is more acute in reconfigurable systems since the compiler must generate the data path and the sequence ...

Keywords: VHDL, compilation, high-level synthesis, reconfigurable computing, reuse analysis

Results 1 - 20 of 42

Result page: [1](#) [2](#) [3](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)